Finding Largest Rectangles in Convex Polygons*

Sergio Cabello[1] Otfried Cheong[2] Christian Knauer[3] Lena Schlipf[4]

October 8, 2014

Abstract

We consider the following geometric optimization problem: find a maximum-area rectangle and a maximum-perimeter rectangle contained in a given convex polygon with n vertices. We give exact algorithms that solve these problems in time $O(n^3)$. We also give $(1-\varepsilon)$-approximation algorithms that take time $O(\varepsilon^{-3/2} + \varepsilon^{-1/2} \log n)$.

Keywords: geometric optimization; approximation algorithm; convex polygon; inscribed rectangle.

arXiv:1405.1223v2 [cs.CG] 6 Oct 2014

1 Introduction

Computing a largest rectangle contained in a polygon (with respect to some appropriate measure) is a well studied problem. Previous results include computing largest *axis-aligned* rectangles, either in convex polygons [2] or simple polygons (possibly with holes) [5], and computing largest *fat* rectangles in simple polygons [6].

Here we study the problem of finding a maximum-area rectangle and a maximum-perimeter rectangle contained in a given convex polygon with n vertices. We give exact $O(n^3)$-time algorithms and $(1-\varepsilon)$-approximation algorithms that take time $O(\varepsilon^{-3/2}+\varepsilon^{-1/2}\log n)$. (For maximizing the perimeter we allow the degenerate solution consisting of a single segment whose perimeter is twice its length.) To the best of our knowledge, apart from a straightforward $\mathcal{O}(n^4)$-time algorithm, there is no other exact algorithm known so far.

Our approximation algorithm to maximize the area improves the previous results by Knauer et al. [7], who give a deterministic $(1-\varepsilon)$-approximation algorithm with running time $O(\varepsilon^{-2}\log n)$ and a Monte Carlo $(1-\varepsilon)$-approximation algorithm with running time $O(\varepsilon^{-1}\log n)$. We are not aware of previous $(1-\varepsilon)$-approximation algorithms to maximize the perimeter.

2 Preliminaries

Notation. We use C for arbitrary convex bodies and P for convex polygons.

Let $\mathcal{U}$ be the set of unit vectors in the plane. For each $u \in \mathcal{U}$ and each convex body C, the *directional width* of C in direction u, denoted by $\mathrm{dwidth}(u, C)$, is the length of the orthogonal projection of C onto any line parallel to u. Thus

$$\mathrm{dwidth}(u, C) \;=\; \max_{p\in C}\langle p, u\rangle - \min_{p\in C}\langle p, u\rangle,$$

where $\langle\cdot,\cdot\rangle$ denotes the scalar product.

*Supported by the Slovenian Research Agency, program P1-0297, projects J1-4106 and L7-5459; by the ESF EuroGIGA project (project GReGAS) of the European Science Foundation; and by NRF grant 2011-0030044 (SRC-GAIA), funded by the government of Korea.

[1]Department of Mathematics, IMFM, and Department of Mathematics, FMF, University of Ljubljana, Slovenia. Part of the work was done while visiting KAIST, Korea.

[2]Department of Computer Science, KAIST, Daejeon, Korea.

[3]Institut für Informatik, Universität Bayreuth, Bayreuth, Germany

[4]Institute of Computer Science, Freie Universität Berlin, Germany.

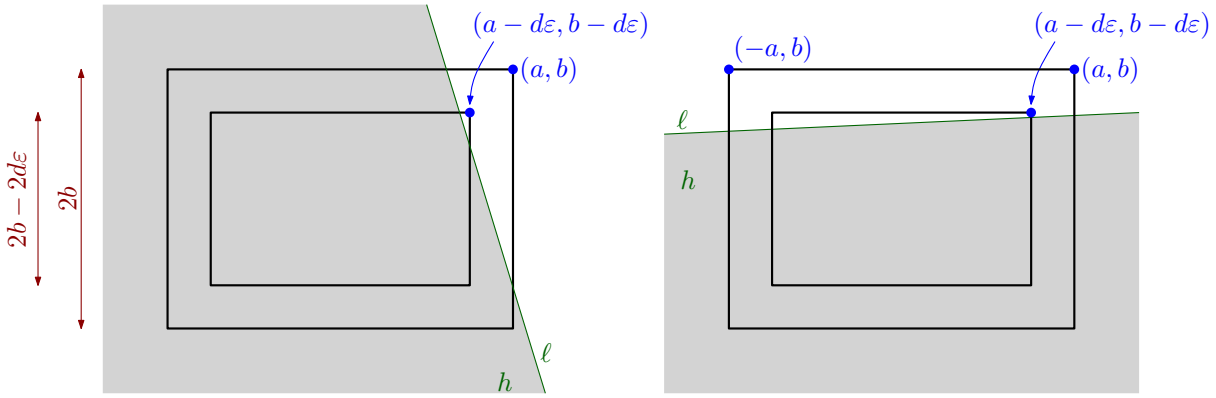

Figure 1: Proof of Lemma 3.

For a convex body C and a parameter $\varepsilon \in (0, 1)$, an *ε-kernel* for C is a convex body $C_\varepsilon \subseteq C$ such that

$$\forall u \in \mathcal{U} : \quad (1 - \varepsilon) \cdot \operatorname{dwidth}(u, C) \leqslant \operatorname{dwidth}(u, C_\varepsilon).$$

The diameter of C is the distance between the two furthest points of C. It is easy to see that it equals

$$\max_{u \in \mathcal{U}} \operatorname{dwidth}(u, C).$$

Ahn et al. [1] show how to compute an ε-kernel. Their algorithm uses the following type of primitive operations for C:

- given a direction $u \in \mathcal{U}$, find an extremal point of C in the direction u;
- given a line ℓ, find $C \cap \ell$.

Let T_C be the time needed to perform each of those primitive operations. We will use T_C as a parameter in some of our running times. When C is a convex n-gon whose boundary is given as a sorted array of vertices or as a binary search tree, we have $T_C = O(\log n)$ [4, 10]. Ahn et al. show the following result.

Lemma 1 (Ahn et al. [1])**.** *Given a convex body C and a parameter $\varepsilon \in (0, 1)$, we can compute in $O(\varepsilon^{-1/2} T_C)$ time an ε-kernel of C with $O(\varepsilon^{-1/2})$ vertices.*

Lemma 2. *Let C_ε be an ε-kernel for C. If φ is an invertible affine mapping, then $\varphi(C_\varepsilon)$ is an ε-kernel for $\varphi(C)$.*

Proof. The ratio of directional widths for convex bodies is invariant under invertible affine transformations. This means that

$$\forall u \in \mathcal{U} : \quad 1 - \varepsilon \leqslant \frac{\operatorname{dwidth}(u, C_\varepsilon)}{\operatorname{dwidth}(u, C)} = \frac{\operatorname{dwidth}(u, \varphi(C_\varepsilon))}{\operatorname{dwidth}(u, \varphi(C))}$$

and thus $\varphi(C_\varepsilon)$ is an ε-kernel for $\varphi(C)$. □

Lemma 3. *Assume that C contains the rectangle $R = [-a, a] \times [-b, b]$, that C has diameter d, and that C_ε is an ε-kernel for C. Then C_ε contains the axis-parallel rectangle $S = [-a+d\varepsilon, a-d\varepsilon] \times [-b+d\varepsilon, b-d\varepsilon]$.*

Proof. The statement is empty if $a < d\varepsilon$ or $b < d\varepsilon$, so assume that $a, b \geqslant d\varepsilon$. For the sake of contradiction, assume that S is not contained in C_ε. This means that one vertex of S is not contained in C_ε. Because of symmetry, we can assume that $s = (a - d\varepsilon, b - d\varepsilon)$ is not contained in C_ε. Since C_ε is convex and $s \notin C_\varepsilon$, there exists a closed halfplane h that contains C_ε but does not contain s. Let ℓ be the boundary of h.

We next argue that R has some vertex at distance at least $d\varepsilon$ from h (and thus ℓ); see Figure 1 for a couple of cases. If ℓ has negative slope and h is its lower halfplane, then the distance from (a, b) to ℓ is at least $d\varepsilon$. If ℓ has negative slope and h is its upper halfplane, then the distance from $(-a, -b)$ to ℓ is at least $2b - d\varepsilon \geqslant d\varepsilon$. If ℓ has positive slope, then $(-a, b)$ or $(a, -b)$ are at distance at least $d\varepsilon$ from h.

Since $R \subseteq C$, for the direction u orthogonal to ℓ we have

$$\mathrm{dwidth}(u, C) - \mathrm{dwidth}(u, C_\varepsilon) > d\varepsilon \geqslant \varepsilon \cdot \mathrm{dwidth}(u, C),$$

where we have used the assumption that $\mathrm{dwidth}(u, C) \leqslant d$. This means that

$$(1 - \varepsilon) \cdot \mathrm{dwidth}(u, C) \;>\; \mathrm{dwidth}(u, C_\varepsilon),$$

which contradicts that C_ε is an ε-kernel for C. □

3 Exact algorithms

Let $e_1, \ldots, e_n$ be the edges of the convex polygon P. For each edge e_i of P, let h_i be the closed halfplane defined by the line supporting e_i that contains P. Since P is convex, we have $P = \bigcap_i h_i$.

We parameterize the set of *parallelograms* in the plane by points in $\mathbb{R}^6$, as follows. We identify each 6-dimensional point $(x_1, x_2, u_1, u_2, v_1, v_2)$ with the triple $(x, u, v) \in (\mathbb{R}^2)^3$, where $x = (x_1, x_2)$, $u = (u_1, u_2)$, and $v = (v_1, v_2)$. The triple $(x, u, v) \in \mathbb{R}^6$ corresponds to the parallelogram $\Diamond(x, u, v)$ with vertices

$$x,\ x + u,\ x + v,\ x + u + v.$$

Thus, x describes a vertex of the parallelogram $\Diamond(x, u, v)$, while u and v are vectors describing the edges of $\Diamond(x, u, v)$. This correspondence is not bijective because, for example,

$$\Diamond(x, u, v) \;=\; \Diamond(x + u + v, -u, -v) \;=\; \Diamond(x, v, u).$$

Nevertheless, each parallelogram is $\Diamond(x, u, v)$ for some $(x, u, v) \in \mathbb{R}^6$: the parallelogram given by the vertices $p_1p_2p_3p_4$ in clockwise (or counterclockwise) order is $\Diamond(p_1, p_2 - p_1, p_4 - p_1)$.

We are interested in the parallelograms contained in P. To this end we define

$$\Pi(P) \;=\; \big\{(x, u, v) \in \mathbb{R}^6 \mid \Diamond(x, u, v) \subseteq P\big\}.$$

Since P is convex, a parallelogram is contained in P if and only if each vertex of the parallelogram is in P. Therefore

$$\begin{aligned}
\Pi(P) \;&=\; \big\{(x, u, v) \in \mathbb{R}^6 \mid x,\ x + u,\ x + v,\ x + u + v \in P\big\} \\
&=\; \big\{(x, u, v) \in \mathbb{R}^6 \mid \forall i : x,\ x + u,\ x + v,\ x + u + v \in h_i\big\} \\
&=\; \bigcap_i \big\{(x, u, v) \in \mathbb{R}^6 \mid x,\ x + u,\ x + v,\ x + u + v \in h_i\big\}.
\end{aligned}$$

Since $\Pi(P)$ is trivially bounded, it follows that $\Pi(P)$ is a convex polytope in $\mathbb{R}^6$ defined by $4n$ linear constraints. The Upper Bound Theorem [9] implies that $\Pi(P)$ has combinatorial complexity at most $O(n^3)$. Chazelle's algorithm [3] gives a triangulation of the boundary of $\Pi(P)$ in $O(n^3)$ time; Seidel [11] calls this the boundary description of a polytope. From the triangulation of the boundary we can construct a triangulation of $\Pi(P)$: chose an arbitrary vertex x of $\Pi(P)$ and add it to each simplex of the triangulation of the boundary of $\Pi(P)$ that does not contain x. (One can also use a point in the interior of $\Pi(P)$.)

The set of *rectangles* is obtained by restricting our attention to triples (x, u, v) with $\langle u, v\rangle = 0$, where $\langle \cdot, \cdot \rangle = 0$ again denotes the scalar product of two vectors. This constraint is non-linear. Because of this, it is more convenient to treat each simplex of a triangulation of $\Pi(P)$ separately. When $\langle u, v\rangle = 0$, the area of $\Diamond(x, u, v)$ is $|u| \cdot |v|$.

Consider any simplex $\triangle$ of the triangulation of $\Pi(P)$. Finding the maximum area rectangle restricted to $\triangle$ corresponds to the problem

$$\begin{aligned}
\mathrm{opt}(\triangle) \;=\; &\max\ |u|^2 \cdot |v|^2 \\
&\text{s.t. } (x, u, v) \in \triangle \\
&\quad \langle u, v\rangle = 0
\end{aligned}$$

This is a constant-size problem. It has 6 variables and a constant number of constraints; all constraints but one are linear. The optimization function has degree four. In any case, each such problem can be solved in constant time. When the problem is not feasible, we set $\mathrm{opt}(\triangle) = 0$.

Taking the best rectangle over all simplices of a triangulation of $\Pi(P)$, we find a maximum area rectangle. Thus, we return $\arg\max_\triangle \mathrm{opt}(\triangle)$. We have shown the following.

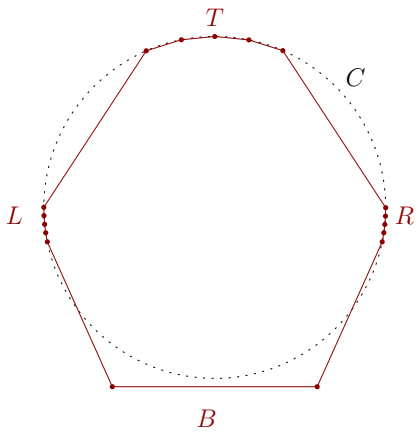

Figure 2: Outline of the construction.

Theorem 4. *Let P be a convex polygon with n vertices. In time $O(n^3)$ we can find a maximum-area rectangle contained in P.*

To maximize the perimeter, we apply the same approach. For each simplex $\triangle$ in a triangulation of $\Pi(P)$ we have to solve the following problem:

$$\begin{aligned} \text{opt}(\triangle) \;=\; & \max\ |u| + |v| \\ & \text{s.t. } (x, u, v) \in \triangle \\ & \langle u, v \rangle = 0 \end{aligned}$$

Combining the solutions over all simplices of the triangulation we obtain the following.

Theorem 5. *Let P be a convex polygon with n vertices. In time $O(n^3)$ we can find a maximum-perimeter rectangle contained in P.*

4 Combinatorially distinct rectangles

The reader may wonder if the algorithm of the previous section cannot be improved: It constructs the space $\Pi(P)$ of all *parallelograms* contained in P, and then considers the intersection with the manifold $\langle u, v \rangle = 0$ corresponding to the rectangles. If the complexity of this intersection was smaller than $\Theta(n^3)$, then we should avoid constructing the entire parallelogram space $\Pi(P)$ first.

In this section we show that this is not the case: the complexity of the space of rectangles that fit inside P, that is, the complexity of the intersection of $\Pi(P)$ with the manifold $\langle u, v \rangle = 0$, is already $\Theta(n^3)$ in the worst case. Therefore, asymptotically we are not loosing anything by considering all parallelograms, instead of directly concentrating on rectangles.

To this end, let us call two rectangles contained in a convex polygon P ***combinatorially distinct*** if their vertices are incident to a different subset of edges of P. We are going to show the following: for every sufficiently large n there is a polygon P with n vertices that contains $\Theta(n^3)$ combinatorially distinct rectangles. This shows that any algorithm iterating over all combinatorially distinct rectangles contained in P needs at least $\Omega(n^3)$ time. Our algorithm falls in this category.

We provide an informal overview of the construction, see Figure 2. For simplicity, we are going to use $3n + 2$ vertices. Consider the circle C of unit radius centered at the origin. We are going to select some points on C, plus two additional points. The polygon P is then described as the convex hull of these points. The points are classified into 4 groups. We have a group L of n points placed densely on the left side of C. We have another group R of n points placed densely on the right side of C. The third group T, also with n points, is placed on the upper part of C. The points of T are more spread out and will be chosen carefully. Finally, we construct a group B of two points placed below C. The construction will have the property that, for any edge e_L defined by L, any edge e_R defined by R, and any edge e_T defined by T, there is a rectangle contained in P with vertices on the edges e_L, e_R and e_T. The points B are needed only to make sure that the bottom part of the rectangle is contained in P; they do not have any particular role.

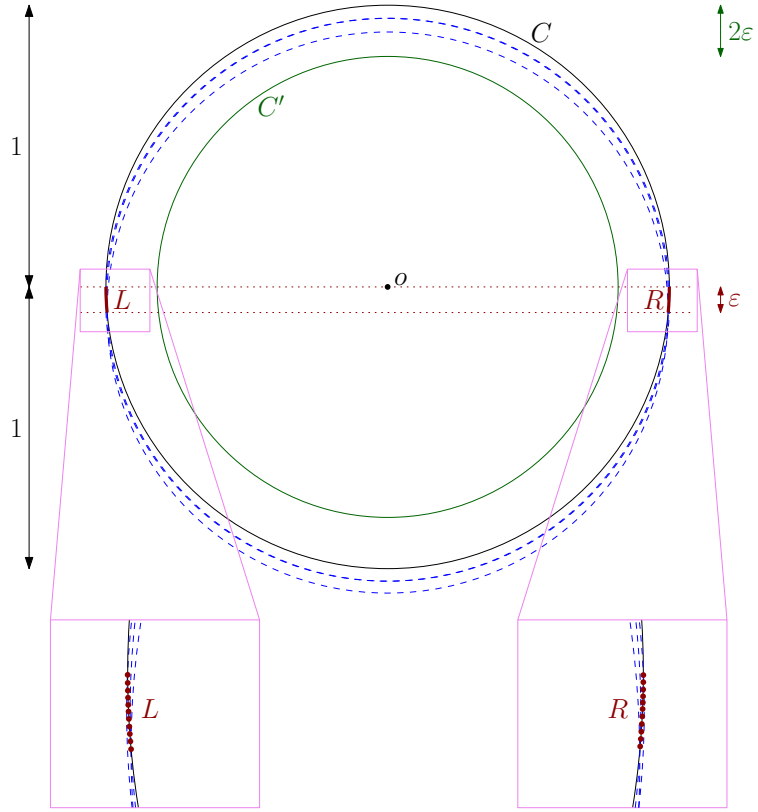

Figure 3: Detail of the construction of L and B. Some circles $C_{r\ell}$ are in dashed blue.

Theorem 6. *For any sufficiently large value of n there is a polygon P with n vertices such that P contains $\Theta(n^3)$ combinatorially distinct rectangles.*

Proof. Let C be the circle of unit radius centered at the origin o, and let C' be the circle of radius $1-2\varepsilon$ centered at o, for a small value $\varepsilon > 0$ to be chosen later on. Let H_ε be the horizontal strip defined by $-\varepsilon \leqslant y \leqslant 0$, see Figure 3. Select a set R of n points in $C \cap H_\varepsilon$ with positive x-coordinate and select a set L of n points in $C \cap H_\varepsilon$ with negative x-coordinate. For every r on a segment connecting consecutive points of R and every ℓ on a segment connecting consecutive points of L, let $C_{r\ell}$ be the circle with diameter $r\ell$. We now observe that the upper semicircle of $C_{r\ell}$ with endpoints r and ℓ lies between C and C'. This follows from the fact that $C_{r\ell}$ has radius at least $1-\varepsilon$ and the center of $C_{r\ell}$ is at most ε apart from o.

Figure 4 should help with the continuation of the construction. We place a set T of n points on the upper side of C. We want the following additional property: for any two consecutive points t ant t' of T, the segment tt' intersects C'. If we select $\varepsilon > 0$ sufficiently small, then C and C' are close enough that we can select the n points needed to construct T. Finally, we choose a set B of two points below C, as shown in Figure 2. The final polygon P is the convex hull of $L \cup R \cup T \cup B$. This finishes the description of the polygon P.

Consider any edge e_R defined by two consecutive vertices of R and chose a point r on e_R. Similarly, consider any edge e_L defined by two consecutive vertices of L and chose a point ℓ on e_L. Let e_T be an edge of P defined by two consecutive points of T. By construction, the circle $C_{r\ell}$ with diameter $r\ell$ intersects the segment e_T in some point, let's call it p. This means that the triangle $\triangle(r,p,\ell)$ has a right angle at p. Let q be the point on $C_{r\ell}$ such that pq is a diameter of $C_{r\ell}$. Then the quadrilateral formed by r, p, ℓ, and q is a rectangle contained in P. Each choice of an edge e_R defined by R, an edge e_L defined by L, and an edge e_T defined by T results in a combinatorially distinct rectangle. Therefore there are $\Omega(n^3)$ combinatorially distinct rectangles contained in P. □

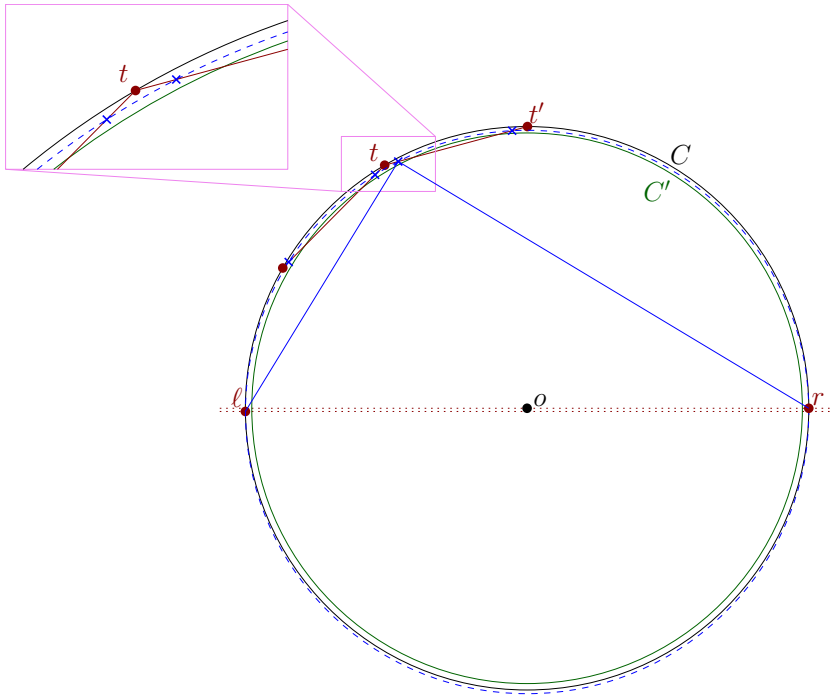

Figure 4: Detail of the construction of T.

5 Approximation algorithm to maximize the area

The algorithm is very simple: we compute an $(\varepsilon/32)$-kernel $C_{\varepsilon/32}$ for the input convex body C, compute a maximum-area rectangle contained in $C_{\varepsilon/32}$ and return it. We next show that this algorithm indeed returns a $(1-\varepsilon)$-approximation.

Let R_{opt} be a maximum-area rectangle contained in C, and let φ be an affine transformation such that $\varphi(R_{\text{opt}})$ is the square $[-1,1]^2$.

Lemma 7. *The diameter of $\varphi(C)$ is at most 16.*

Proof. We will show that $\varphi(C)$ is contained in the disk centered at the origin $o = (0,0)$ of radius 8, which implies the result.

Any convex body contains a rectangle with at least half of its area [8]. Therefore $\operatorname{area}(R_{\text{opt}})/\operatorname{area}(C) \geqslant 1/2$.

Any invertible affine transformation does not change the ratio between areas of objects. Therefore

$$\frac{1}{2} \leqslant \frac{\operatorname{area}(R_{\text{opt}})}{\operatorname{area}(C)} = \frac{\operatorname{area}(\varphi(R_{\text{opt}}))}{\operatorname{area}(\varphi(C))} = \frac{4}{\operatorname{area}(\varphi(C))}$$

and thus $\operatorname{area}(\varphi(C)) \leqslant 8$.

Assume, for the sake of reaching a contradiction, that $\varphi(C)$ has a point p at distance larger than 8 from the origin o. See Figure 5. Let s be the line segment of length 2 centered at the origin o and orthogonal to the segment op. Since s is contained in the square $[-1,1]^2$, it is contained in $\varphi(C)$. Therefore $\varphi(C)$ contains the convex hull of $s \cup \{p\}$, which is a triangle of area larger than 8, and we get a contradiction. It follows that $\varphi(C)$ is contained in a disk centered at the origin 0 of radius 8. □

Lemma 8. *Let C_ε be an ε-kernel for C. Then C_ε contains a rectangle with area at least $(1-32\varepsilon)\cdot \operatorname{area}(R_{\text{opt}})$.*

Proof. Because of Lemma 2, $\varphi(C_\varepsilon)$ is an ε-kernel for $\varphi(C)$. Since $\varphi(C)$ contains $[-1,1]^2$ and has diameter at most 16 due to Lemma 7, Lemma 3 with $a = b = 1$ implies that $\varphi(C_\varepsilon)$ contains the square $S = [-t,t]^2$, where $t = 1 - 16\varepsilon$.

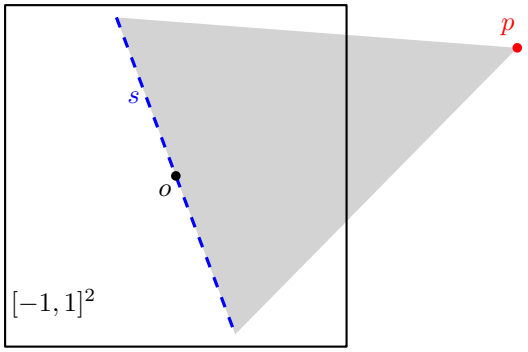

Figure 5: Triangle considered in the proof of Lemma 7. The point p is drawn closer to the origin o than it is assumed.

Since S is obtained by scaling $[-1,1]^2 = \varphi(R_{\text{opt}})$ by $1-16\varepsilon$, its preimage $R = \varphi^{-1}(S)$ is obtained by scaling R_{opt} by $1-16\varepsilon$ about its center. It follows that R is a rectangle with area

$$\text{area}(R) \;=\; (1-16\varepsilon)^2 \cdot \text{area}(R_{\text{opt}}) \;\geqslant\; (1-32\varepsilon) \cdot \text{area}(R_{\text{opt}}),$$

and the lemma follows. □

Theorem 9. *Let C be a convex body in the plane. For any given $\varepsilon \in (0,1)$, we can find a $(1-\varepsilon)$-approximation to the maximum-area rectangle contained in C in time $O(\varepsilon^{-1/2}T_C + \varepsilon^{-3/2})$.*

Proof. First, we compute an $(\varepsilon/32)$-kernel $C_{\varepsilon/32}$ to C. We then compute a maximum-area rectangle contained in $C_{\varepsilon/32}$ and return it. This finishes the description of the algorithm.

Because of Lemma 8, $C_{\varepsilon/32}$ contains a rectangle of area at least $(1-\varepsilon)\cdot\text{area}(R_{\text{opt}})$, where R_{opt} is a maximum-area rectangle contained in C. Therefore, the algorithm returns a $(1-\varepsilon)$-approximation to the maximum-area rectangle.

Computing $C_{\varepsilon/32}$ takes time $O((\varepsilon/32)^{-1/2}T_C) = O(\varepsilon^{-1/2}T_C)$ because of Lemma 1. Since $C_{\varepsilon/32}$ has $O((\varepsilon/32)^{-1/2}) = O(\varepsilon^{-1/2})$ vertices, finding a largest rectangle contained in $C_{\varepsilon/32}$ takes time $O(\varepsilon^{-3/2})$ because of Theorem 4. □

Corollary 10. *Let C be a convex polygon with n vertices given as a sorted array or a balanced binary search tree. For any given $\varepsilon \in (0,1)$, we can find a $(1-\varepsilon)$-approximation to the maximum-area rectangle contained in C in time $O(\varepsilon^{-1/2}\log n + \varepsilon^{-3/2})$.*

Proof. In this case $T_C = O(\log n)$. □

6 Approximation algorithm to maximize the perimeter

The approximation algorithm is the following: we compute an $(\varepsilon/16)$-kernel $C_{\varepsilon/16}$ for the input convex body C, compute a maximum-perimeter rectangle $R_{\varepsilon/16}$ contained in $C_{\varepsilon/16}$, and return it. We next show that this indeed computes a $(1-\varepsilon)$-approximation.

Since the algorithm is independent of the coordinate axes, we can assume that the maximum-perimeter rectangle contained in C is an axis-parallel rectangle $R_{\text{opt}} = [-a,a]\times[-b,b]$ with $b \leqslant a$. We distinguish two cases depending on the aspect ratio $b/a \leqslant 1$ of R_{opt}. When $b/a \leqslant \varepsilon/2$, then the longest segment contained in C is a good approximation to R_{opt}. When $b/a > \varepsilon/2$, then R_{opt} is fat enough that we can use Lemma 3 to obtain a large-perimeter rectangle.

Lemma 11. *We have* $\text{peri}(R_{\varepsilon/16}) \geqslant (1-\varepsilon)\cdot\text{peri}(R_{\text{opt}})$.

Proof. We assume first that $b \leqslant (\varepsilon/2)a$. This implies $\text{peri}(R_{\text{opt}}) \leqslant 4(1+\varepsilon/2)a$. Because C contains R_{opt}, the directional width of C in the horizontal direction is at least $2a$. The diameter of $C_{\varepsilon/16}$ is therefore at least $(1-\varepsilon/16)2a$, and so $C_{\varepsilon/16}$ contains a segment of perimeter at least $4(1-\varepsilon/16)a$. The lemma then follows from

$$4(1-\varepsilon/16)a \geqslant (1-\varepsilon)\cdot 4(1+\varepsilon/2)a.$$

It remains to consider the case $b > (\varepsilon/2)a$. Let D be the disk of radius $2a$ centered at the origin. We have $C \subset D$, as otherwise C contains a segment of perimeter strictly larger than $4a \geqslant 2a + 2b = \operatorname{peri}(R_{\text{opt}})$, contradicting the optimality of R_{opt}. It follows that the diameter of C is at most $4a$.

By Lemma 3, $C_{\varepsilon/16}$ contains the axis-parallel rectangle $S = [-a + t, a - t] \times [-b + t, b - t]$, where $t = a\varepsilon/4$. We have

$$\operatorname{peri}(S) = \operatorname{peri}(R_{\text{opt}}) - 8t = \operatorname{peri}(R_{\text{opt}}) - 2a\varepsilon > (1 - \varepsilon) \operatorname{peri}(R_{\text{opt}}). \qquad \square$$

Theorem 12. *Let C be a convex body in the plane. For any given $\varepsilon \in (0, 1)$, we can find a $(1 - \varepsilon)$-approximation to the maximum-perimeter rectangle contained in C in time $O(\varepsilon^{-1/2} T_C + \varepsilon^{-3/2})$.*

Proof. First, we compute an $(\varepsilon/16)$-kernel $C_{\varepsilon/16}$ for C. We then compute a maximum-perimeter rectangle contained in $C_{\varepsilon/16}$ and return it. This finishes the description of the algorithm.

By Lemma 11 $C_{\varepsilon/16}$ contains a rectangle of perimeter at least $(1 - \varepsilon) \cdot \operatorname{peri}(R_{\text{opt}})$, where R_{opt} is a maximum-perimeter rectangle contained in C. Therefore, the algorithm returns a $(1 - \varepsilon)$-approximation to the maximum-perimeter rectangle.

Computing $C_{\varepsilon/16}$ takes time $O(\varepsilon^{-1/2} T_C)$ because of Lemma 1. Finding the maximum-perimeter rectangle contained in $C_{\varepsilon/16}$ takes time $O(\varepsilon^{-3/2})$ because of Theorem 5. $\square$

Corollary 13. *Let C be a convex polygon with n vertices given as a sorted array or a balanced binary search tree. For any given $\varepsilon \in (0, 1)$, we can find a $(1 - \varepsilon)$-approximation to the maximum-perimeter rectangle contained in C in time $O(\varepsilon^{-1/2} \log n + \varepsilon^{-3/2})$.*

Proof. In this case $T_C = O(\log n)$. $\square$

References

[1] H.-K. Ahn, P. Braß, O. Cheong, H.-S. Na, C.-S. Shin, and A. Vigneron. Inscribing an axially symmetric polygon and other approximation algorithms for planar convex sets. *Comput. Geom.*, 33(3):152–164, 2006.

[2] H. Alt, D. Hsu, and J. Snoeyink. Computing the largest inscribed isothetic rectangle. In *Proc. CCCG'95*, pages 67–72, 1995.

[3] B. Chazelle. An optimal convex hull algorithm in any fixed dimension. *Discrete & Computational Geometry*, 10:377–409, 1993.

[4] B. Chazelle and D. P. Dobkin. Intersection of convex objects in two and three dimensions. *J. ACM*, 34(1):1–27, 1987.

[5] K. Daniels, V. Milenkovic, and D. Roth. Finding the largest area axis-parallel rectangle in a polygon. *Computational Geometry*, 7(12):125–148, 1997.

[6] O. A. Hall-Holt, M. J. Katz, P. Kumar, J. S. B. Mitchell, and A. Sityon. Finding large sticks and potatoes in polygons. In *Proc. SODA'06*, pages 474–483, 2006.

[7] C. Knauer, L. Schlipf, J. M. Schmidt, and H. R. Tiwary. Largest inscribed rectangles in convex polygons. *J. Discrete Algorithms*, 13:78–85, 2012.

[8] M. Lassak. Approximation of convex bodies by rectangles. *Geometriae Dedicata*, 47(1):111–117, 1993.

[9] P. McMullen. The maximum numbers of faces of a convex polytope. *Mathematika*, 17:179–184, 1971.

[10] M. Reichling. On the detection of a common intersection of k convex polyhedra. In *Computational Geometry and its Applications*, volume 333 of *LNCS*, pages 180–186. Springer Berlin Heidelberg, 1988.

[11] R. Seidel. Convex hull computations. In J. E. Goodman and J. O'Rourke, editors, *Handbook of Discrete and Computational Geometry*, pages 361–375. CRC Press, Inc., Boca Raton, FL, USA, 1997.